

The IMP game: Learnability, approximability and adversarial learning beyond Σ_1^0

Michael Brand^a, David L. Dowe^a

^a*Faculty of IT (Clayton), Monash University, Clayton, VIC 3800, Australia*

Abstract

We introduce a problem set-up we call the Iterated Matching Pennies (IMP) game and show that it is a powerful framework for the study of three problems: adversarial learnability, conventional (i.e., non-adversarial) learnability and approximability. Using it, we are able to derive the following theorems. (1) It is possible to learn by example all of $\Sigma_1^0 \cup \Pi_1^0$ as well as some supersets; (2) in adversarial learning (which we describe as a pursuit-evasion game), the pursuer has a winning strategy (in other words, Σ_1^0 can be learned adversarially, but Π_1^0 not); (3) some languages in Π_1^0 cannot be approximated by any language in Σ_1^0 .

We show corresponding results also for Σ_i^0 and Π_i^0 for arbitrary i .

Keywords: Turing machine, recursively enumerable, decidable, approximation, matching pennies, halting, halting problem, elusive model paradox, red herring sequence, learnability, Nash equilibrium, approximability, adversarial learning

1. Introduction

This paper deals with three widely-discussed topics: approximability, conventional learnability and adversarial learnability, and introduces a unified framework in which all three can be studied.

First, consider approximability. Turing's seminal 1936 result [21] demonstrated that some languages that can be accepted by Turing machines (TMs) are not decidable. Otherwise stated, some R.E. languages are not recursive. Equivalently: some co-R.E. languages are not R.E.; any R.E. language must differ from them by at least one word. However, the diagonalisation process by which this result was originally derived makes no stronger claim regarding the number of words differentiating a co-R.E. language and an R.E. one. It merely shows one example of a word where a difference must exist.

We extend this original result by showing that some co-R.E. languages are, in some sense, as different from any R.E. language as it is possible to be.

To formalise this statement, consider an arbitrary (computable) enumeration, w_1, w_2, \dots , over the complete language (the language that includes all words over the chosen alpha-

Email addresses: michael.brand@monash.edu (Michael Brand), david.dowe@monash.edu (David L. Dowe)

bet). Over this enumeration, $\{w_i\}$, we define a distance metric, **dissimilarity**, between two languages, L_1 and L_2 , as follows.

$$\text{DisSim}(L_1, L_2) \equiv \limsup_{n \rightarrow \infty} \frac{|(L_1 \triangle L_2) \cap \{w_1, \dots, w_n\}|}{n},$$

where $L_1 \triangle L_2$ is the symmetric difference. We note that the value of $\text{DisSim}(L_1, L_2)$ depends on the enumeration chosen, and therefore, technically, $\text{DisSim}(\cdot) = \text{DisSim}_{\{w_i\}}(\cdot)$. However, all results in this paper are true for all possible choices of the enumeration, for which reason we omit the choice of enumeration, opting for this more simplified notation.

$\text{DisSim}(L_1, L_2)$ ranges between 0 (the languages are essentially identical) and 1 (the languages are completely dissimilar).

We prove:

Theorem 1. *There is a co-R.E. language \bar{L} such that every R.E. language has a dissimilarity distance of 1 from \bar{L} .*

Consider now learnability. Learnability is an important concept in statistics, econometrics, machine learning, inductive inference, data mining and other fields. This has been discussed by E. M. Gold and by L. G. Valiant in terms of language identification in the limit [7, 22], and also in statistics via the notion of statistical consistency, also known as “completeness” (converging arbitrarily closely in the limit to an underlying true model).

Following upon his convergence results in [17], Solomonoff writes [20, sec. 2 (Completeness and Incomputability)]:

“It is notable that completeness and incomputability are complementary properties: It is easy to prove that any complete prediction method must be incomputable. Moreover, any computable prediction method can not be complete – there will always be a large space of regularities for which its predictions are catastrophically poor.”

In other words, in Solomonoff’s problem set-up it is impossible for a Turing machine to learn every R.E. language: every computable learner is limited.

Nevertheless, in the somewhat different context within which we study learnability, we are able to show that this tension does not exist: a Turing machine can learn any computable language. Moreover, we will consider a set of languages that includes, as a proper subset of it, the languages $\Sigma_1^0 \cup \Pi_1^0$ and will prove that while no deterministic learning algorithm can learn every language in the set, a probabilistic one can (with probability 1), and a mixed strategy involving several deterministic learning algorithms can approximate this arbitrarily well.¹

Lastly, consider adversarial learning [12, 11, 9]. This is different from the conventional learning scenario described above in that while in conventional learning we attempt to

¹Here and elsewhere we use the standard notations for language families in the arithmetical hierarchy [15]: Σ_1^0 is the set of recursively enumerable languages, Π_1^0 is the set of co-R.E. languages.

converge to an underlying “true model” based on given observations, adversarial learning is a multi-player process in which each participant can observe (to some extent) other players’ predictions and adjust their own actions accordingly. This game-theoretic set-up becomes of practical importance in many scenarios. For example, in online bidding bidders use information available to them (e.g., whether they won a particular auction) to learn the strategy used by competing bidders, so as to be able to optimise their own strategy accordingly.

We consider, specifically, an adversarial learning scenario in which one player (the pursuer) attempts to copy a second player, while the second player (the evader) is attempting to avoid being copied. Specifically, each player generates a bit (0 or 1) and the pursuer wins if the two bits are equal while the evader wins if they are not. Though on the face of it this scenario may seem symmetric, we show that the pursuer has a winning strategy.

To attain all these results (as well as their higher-Turing-degree equivalents), we introduce a unified framework in which these questions and related ones can all be studied. The set-up used is an adaptation of one initially introduced by Scriven [16] of a predictor and a contrapredictive (or avoider) effectively playing what we might nowadays describe as a game of iterated matching pennies. In Section 2, we give a formal description of this problem set-up and briefly describe its historical evolution. In Section 3, we explain the relevance of the set-up to the learnability and approximability problems and analyse, as an example case, adversarial learning in the class of decidable languages. In Section 4, we extend the analysis to adversarial learning in all other classes in the arithmetical hierarchy, and in particular to Turing machines.

In Sections 5 and 6 we then return to conventional learnability and to approximability, respectively, and prove the remaining results by use of the set-up developed, showing how it can be adapted to these problems.

2. Matching Pennies

The matching pennies game is a zero-sum two-player game where each player is required to output a bit. If the two bits are equal, this is a win for Player “=”; if they differ, this is a win for Player “≠”. The game is a classic example used in teaching mixed strategies [see, e.g. 6, pp. 283–284]: its only Nash equilibrium [14, 13] is a mixed strategy wherein each player chooses each of the two options with probability 1/2.

Consider, now, an iterative version of this game, where at each round the players choose a new bit with perfect information of all previous rounds. Here, too, the best strategy is to choose at each round a new bit with probability 1/2 for each option, and with the added caveat that each bit must be independent of all previous bits. In the iterative variation, we define the payoff (of the entire game) to be

$$S = S_{\neq} = \left(\liminf_{N \rightarrow \infty} \sum_{n=1}^N \frac{\delta_n}{2N} \right) + \left(\limsup_{N \rightarrow \infty} \sum_{n=1}^N \frac{\delta_n}{2N} \right) \quad (1)$$

for Player “≠”, where δ_n is 0 if the bits output in the n ’th round are equal and 1 if they

are different. The payoff for Player “=” is

$$S_{=} = 1 - S_{\neq} = \left(\liminf_{N \rightarrow \infty} \sum_{n=1}^N \frac{1 - \delta_n}{2N} \right) + \left(\limsup_{N \rightarrow \infty} \sum_{n=1}^N \frac{1 - \delta_n}{2N} \right) \quad (2)$$

These payoff functions were designed to satisfy the following criteria:

- They are always defined.
- The game is zero-sum and strategically symmetric, except for the essential distinction between a player aiming to copy (Player “=”, the pursuer) and a player aiming for dissimilarity (Player “ \neq ”, the evader).
- The payoff is a function solely of the $\{\delta_i\}$ sequence. (This is important because in the actual IMP game being constructed players will only have visibility into past δ_i , not full information regarding the game’s evolution.)
- Where a limit exists (in the lim sense) to the percentage of rounds to be won by a player, the payoff is this percentage.

In particular, note that when the payoff functions take the value 0 or 1, there exists a limit (in the lim sense) to the percentage of rounds to be won by a player, and in this case the payoff is this limit.

In the case of the strategy pair described above, for example, where bits are determined by independent, uniform-distribution coin tosses, the limit exists and the payoff is 1/2 for both players, indicating that the game is not biased towards either. This is a Nash equilibrium of the game: neither player can ensure a higher payoff for herself as long as the other persists in the equilibrium strategy. The game has other Nash equilibria, but all share the (1/2, 1/2) payoffs.

Above, we describe the players in the game as agents capable of randomisation: they choose a random bit at each new round. However, the game can be played, with the same strategies, also by deterministic agents. For this, consider every possible infinite bit-string as a possible strategy for each of the players. In this case, the game’s Nash equilibrium would be a strategy pair where each player allots a bit-string from a uniform distribution among all options.

We formalise this deterministic outlook on the matching pennies game as follows.

Definition 1 (Iterative Matching Pennies game). An *Iterative Matching Pennies game* (or *IMP*), denoted $\text{IMP}(\Sigma_{=}, \Sigma_{\neq})$, is a two player game where each player chooses a language: Player “=” chooses $L_{=} \in \Sigma_{=}$ and Player “ \neq ” chooses $L_{\neq} \in \Sigma_{\neq}$, where $\Sigma_{=}$ and Σ_{\neq} are two collections of languages over the binary alphabet.

Where $\Sigma_{=} = \Sigma_{\neq} (= \Sigma)$, we denote the game $\text{IMP}(\Sigma)$.

Define Δ_0 to be the empty string and define, for every natural i ,

$$\delta_i \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } \Delta_{i-1} \in L_{=} \Delta L_{\neq} \\ 0 & \text{if } \Delta_{i-1} \notin L_{=} \Delta L_{\neq} \end{cases},$$

$$\Delta_i \stackrel{\text{def}}{=} \Delta_{i-1}\delta_i,$$

Then the payoffs $S_+ = S_+(L_+, L_-)$ and $S_- = S_-(L_+, L_-)$ are as defined in (2) and (1), respectively. The notation “ $\Delta_{i-1}\delta_i$ ” indicates string concatenation.

Player (mixed) strategies in this game are described as distributions, D_+ and D_- , over Σ_+ and Σ_- , respectively. In this case, we define

$$S_+(D_+, D_-) = E(S_+(L_+, L_-)) \quad L_+ \sim D_+, L_- \sim D_-.$$

$$S_-(D_+, D_-) = E(S_-(L_+, L_-)) \quad L_+ \sim D_+, L_- \sim D_-.$$

Note again that the game is zero sum: any pair of strategies, pure or mixed, satisfies

$$S_+(D_+, D_-) + S_-(D_+, D_-) = 1. \quad (3)$$

To better illustrate the dynamics embodied by Definition 1, let us add two more definitions: let

$$O_+(i) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } \Delta_{i-1} \in L_+ \\ 0 & \text{if } \Delta_{i-1} \notin L_+ \end{cases} \quad (4)$$

and let

$$O_-(i) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } \Delta_{i-1} \in L_- \\ 0 & \text{if } \Delta_{i-1} \notin L_- \end{cases}, \quad (5)$$

noting that by Definition 1, $\delta_i = O_+(i) \oplus O_-(i)$, where “ \oplus ” denotes the exclusive or (“xor”) function.

The scenario encapsulated by the IMP game is that of a competition between two players, Player “=” and Player “ \neq ”, where the strategy of the players is encoded in the form of the languages L_+ and L_- , respectively (or distributions over these in the case of mixed strategies).

After i rounds, each player has visibility to the set of results so far. This is encoded by means of Δ_i , a word composed of the characters $\delta_1, \dots, \delta_i$, where each δ_k is 0 if the bits that were output by the two players in round k are equal and 1 if they are not. It is based on this history that the players now generate a new bit: Player “=” generates $O_+(i+1)$ and Player “ \neq ” generates $O_-(i+1)$. The players’ strategies are therefore functions from a word (Δ_i) to a bit ($O_+(i+1)$ for Player “=”, $O_-(i+1)$ for Player “ \neq ”). To encode these strategies in the most general form, we use languages: L_+ and L_- are simply sets containing all the words to which the response is “1”. Our choice of how weak or how strong a player can be is then ultimately in the question of what language family, Σ , its strategy is chosen from.

Once $O_+(i+1)$ and $O_-(i+1)$ are determined, δ_{i+1} is simply their xor (1 if the bits differ, 0 if they are the same), and in this way the definition generates the infinite list of δ_i that is ultimately used to compute the game’s overall payoff for each player.

Were we to actually try and run a real-world IMP competition by directly implementing the definitions above, and were we to try to implement the Nash equilibrium player strategies, we would immediately run into two elements in the set-up that are uncomputable:

first, the choice of a uniform infinitely-long bit-string, our chosen distribution among the potential strategies, is uncomputable (it is a choice among uncountably many elements); second, for a deterministic player (an agent) to output all the bits of an arbitrary (i.e., general) bit-string, that player cannot be a Turing machine. There are only countably many Turing machines, so only countably many bit-strings that can thus be output.

In this paper, we examine the IMP game with several choices for $\Sigma_ =$ and $\Sigma_ \neq$. The main case studied is where $\Sigma_ = \Sigma_ \neq = \Sigma_1^0$. In this case, we still allow player mixed strategies to be uncomputable distributions, but any $L_ =$ and $L_ \neq$ are computable by TMs.

The set-up described here, where Iterated Matching Pennies is essentially described as a pursuit-evasion game, was initially introduced informally by Scriven [16] in order to prove that unpredictability is innate to humans. Lewis and Richardson [10], without explicitly mentioning Turing machines or any (equivalent) models of computation, reinvestigated the model and used it to refute Scriven’s claim, with a proof that hinges on the halting problem, but references it only implicitly.

The set-up was redeveloped independently by Dowe, first in the context of the avoider trying to choose the next number in an integer sequence to be larger (by one) than the (otherwise) best inference that one might expect [1, sec. 0.2.7, p. 545, col. 2 and footnote 211], and then, as in [16], in the context of predicting bits in a sequence [2, p. 455][4, pp. 16–17]. Dowe was the first to introduce the terminology of TMs into the set-up. His aim was to illicit a paradox, which he dubbed “the elusive model paradox”, whose resolution relies on the undecidability of the halting problem. Thus, it would provide an alternative to the method of [21] to prove this undecidability. Variants of the elusive model paradox and of the “red herring sequence” (the optimal sequence to be used by an avoider) are discussed in [3, sec. 7.5], with the paradox also mentioned in [5, sec. 2.2][8, footnote 9].

Yet a third independent incarnation of the model was by Solomonoff, who discussed variants of the elusive model paradox and the red herring sequence in [19, Appendix B] and [18, sec. 3].

We note that the more formal investigations of Dowe and of Solomonoff were in contexts in which the “game” character of the set-up was not explored. Rather, the set-up was effectively a one-player game, where regardless of the player’s choice of next bit, the red herring sequence’s next bit was its reverse. We, on the other hand, return to the original spirit of Scriven’s formulation, investigating the dynamics of the two player game, but do so in a formal setting.

Specifically, we investigate the question of which of the two players (if either) has an advantage in this game, and, in particular, we will be interested in the game’s Nash equilibria, which are the pairs of strategies $(D_ =^*, D_ \neq^*)$ for which

$$S_ =(D_ =^*, D_ \neq^*) = \sup_{D_ =} S_ =(D_ =, D_ \neq^*)$$

and

$$S_ \neq(D_ =^*, D_ \neq^*) = \sup_{D_ \neq} S_ \neq(D_ =^*, D_ \neq).$$

We define

$$\minmax(\Sigma_=_ , \Sigma_{\neq}) = \inf_{D=_} \sup_{D_{\neq}} S(D_=_ , D_{\neq})$$

and

$$\maxmin(\Sigma_=_ , \Sigma_{\neq}) = \sup_{D_{\neq}} \inf_{D=_} S(D_=_ , D_{\neq}),$$

where $D_=_$ is a (potentially incomputable) distribution over $\Sigma_=_$ and D_{\neq} is a (potentially incomputable) distribution over Σ_{\neq} . Where $\Sigma_=_ = \Sigma_{\neq} (= \Sigma)$, we will abbreviate this to $\minmax(\Sigma)$ and $\maxmin(\Sigma)$.

A Nash equilibrium $(D^*_=_ , D^*_{\neq})$ must satisfy

$$S(D^*_=_ , D^*_{\neq}) = \maxmin(\Sigma_=_ , \Sigma_{\neq}) = \minmax(\Sigma_=_ , \Sigma_{\neq}), \quad (6)$$

where, as before, $S = S_{\neq}$.

We note that while it may seem, at first glance, that the introduction of game dynamics into the problems of learnability and approximability inserts an unnecessary complication into their analysis, in fact, we will show that the ability to learn and/or approximate languages, when worded formally, involves a large number of interlocking “lim”, “sup”, “inf”, “lim sup” and “lim inf” clauses that are most naturally expressed in terms of minmax and maxmin solutions, Nash equilibria and mixed strategies.

3. Halting Turing machines

The IMP game serves as a natural platform for investigating adversarial learning: each of the players has the opportunity to learn from all previous rounds, extrapolate from this to the question of what algorithm their adversary is employing and then choose their own course of action to best counteract the adversary’s methods.

Furthermore, where $\Sigma_=_ = \Sigma_{\neq} (= \Sigma)$, IMP serves as a natural arena to differentiate between the learning of a language (e.g., one selected from R.E.) and its complement (e.g., a language selected from co-R.E.), because Player “=”, the copying player, is essentially trying to learn a language from Σ , namely that chosen by Player “ \neq ”, whereas Player “ \neq ” is attempting to learn a language from co- Σ , namely the complement to that chosen by Player “=”. Any advantage to Player “=” can be attributed solely to the difficulty to learn co- Σ by an algorithm from Σ , as opposed to the ability to learn Σ .

To exemplify IMP analysis, consider first the game where $\Sigma = \Delta_1^0$, the set of decidable languages. Because decidable languages are a set known to be closed under complement, we expect Player “ \neq ” to be equally as successful as Player “=” in this variation. Consider, therefore, what would be the Nash equilibria in this case.

Theorem 2. *Let Σ be the set of decidable languages over $\{0,1\}^*$. The game $IMP(\Sigma)$ does not have any Nash equilibria.*

We remark here that most familiar and typically-studied games belong to a family of games where the space of mixed strategies is compact and convex, such as those having

a finite number of pure strategies, and such games necessarily have at least one Nash equilibrium. However, the same is not true for arbitrary games. (For example, the game of “guess the highest number” does not have a Nash equilibrium.) IMP, specifically, does not belong to a game family that guarantees the existence of Nash equilibria.

Proof. We begin by showing that for any (mixed) strategy D_{\neq} ,

$$\sup_{D_{=}} S_{=}(D_{=}, D_{\neq}) = 1. \quad (7)$$

Let $\mathcal{T}_0, \mathcal{T}_1, \dots$ be any (necessarily incomputable) enumeration over those Turing machines that halt on every input, and let L_0, L_1, \dots be the sequence of languages that is accepted by them. The sequence $\{L_i\}$ enumerates (with repetitions) over all languages in $\Sigma = \Delta_1^0$. Under this enumeration we have

$$\lim_{X \rightarrow \infty} \text{Prob}(\exists x \leq X, \text{ such that } L_{\neq} = L_x) = 1; \quad L_{\neq} \sim D_{\neq}.$$

For this reason, for any ϵ there exists an X such that

$$\text{Prob}(\exists x \leq X, \text{ such that } L_{\neq} = L_x) \geq 1 - \epsilon; \quad L_{\neq} \sim D_{\neq}.$$

We devise a strategy, $D_{=}$, to be used by Player $=$. This strategy will be pure: the player will always choose language $L_{=}$, which we will now describe. The language $L_{=}$ is the one accepted by Algorithm 1.

Algorithm 1 Algorithm for learning a mixed strategy

```

1: function CALCULATE BIT( $\Delta$ )
2:    $d \leftarrow \|\Delta\|_1$ . ▷ Number of prediction errors so far.
3:   if  $d > X$  then
4:     Accept.
5:   else if  $\Delta \in L_d$  then
6:     Accept.
7:   else
8:     Reject.
9:   end if
10: end function

```

Note that while the enumeration $\mathcal{T}_0, \mathcal{T}_1, \dots$ is not computable, Algorithm 1 only requires $\mathcal{T}_0, \dots, \mathcal{T}_X$ to be accessible to it, and this can be done because any such finite set of TMs can be hard coded into Algorithm 1.

Consider the game, on the assumption that Player “ \neq ”’s strategy is L_x for $x \leq X$. After at most x prediction errors, Algorithm 1 will begin mimicking a strategy equivalent to L_x and will win every round from that point on.

We see, therefore, that for any $x \in \{0, \dots, X\}$ we have $S_=(L_-, L_x) = 1$, from which we conclude that $S_=(D_-, D_\neq) \geq 1 - \epsilon$ (or, equivalently, $S_\neq(D_-, D_\neq) \leq \epsilon$), in turn proving that for any Nash equilibrium (D^*, D_\neq^*) we necessarily must have

$$\maxmin(\Sigma) = 0. \quad (8)$$

For exactly the symmetric reasons, when $\Sigma = \Delta_1^0$ we also have

$$\minmax(\Sigma) = 1 : \quad (9)$$

Player “ \neq ” can follow a strategy identical to that described in Algorithm 1, except reversing the condition in Step 5.

Because we now have that $\minmax(\Sigma) \neq \maxmin(\Sigma)$, we know that Equation (6) cannot be satisfied for any strategy pair. In particular, there are no Nash equilibria. \square

This result is not restricted to $\Sigma = \Delta_1^0$, the decidable languages, but also to any set of languages that is powerful enough to encode Algorithm 1 and its complement. It is true, for example, for Δ_0^0 as well as for Δ_1^0 with any set of Oracles, i.e., specifically, for any Δ_i^0 .

Definition 2. We say that a collection of languages Σ_\neq is **adversarially learnable** by a collection of strategies $\Sigma_=-$ if $\minmax(\Sigma_-, \Sigma_\neq) = 0$.

If a collection is adversarially learnable by Σ_1^0 , we simply say that it is **adversarially learnable**.

Corollary 2.1. $\forall i, \Delta_i^0$ is not adversarially learnable by Δ_i^0 .

Proof. As was shown in the proof of Theorem 2, $\minmax(\Delta_i^0, \Delta_i^0) = 1$. \square

We proceed, therefore, to the question of how well each player fares when Σ includes non-decidable R.E. languages, and is therefore no longer closed under complement.

4. Adversarial learning

We claim that R.E. languages are adversarially learnable, and that it is therefore not possible to learn the complement of R.E. languages in general, in the adversarial learning scenario.

Theorem 3. The game $IMP(\Sigma_1^0)$ has a strategy, $L_=-$, for Player “ $=$ ” that guarantees $S_\neq(L_-, L_\neq) = 0$ for all L_\neq (and, consequently, also for all distributions among potential L_\neq candidates).

In particular, Σ_1^0 is adversarially learnable.

Proof. We describe $L_=-$ explicitly by means of an algorithm accepting it. This is given in Algorithm 2.

Note that Algorithm 2 does not have any “Accept” or “Reject” statements. It returns a bit only if \mathcal{T}_d returns a bit and does not terminate if \mathcal{T}_d fails to terminate. To actually

Algorithm 2 Algorithm for learning an R.E. language

```
1: function CALCULATE BIT( $\Delta$ )  
2:   Let  $\mathcal{T}_0, \mathcal{T}_1, \dots$  be an enumeration over all Turing machines.  
3:    $d \leftarrow \|\Delta\|_1$ . ▷ Number of prediction errors so far.  
4:   Simulate  $\mathcal{T}_d$   
5: end function
```

simulate \mathcal{T}_d and to encode the enumeration \mathcal{T}_0, \dots , Algorithm 2 can simply use a universal Turing machine, \mathcal{U} , and define the enumeration in a way such that \mathcal{U} accepts the input “ $d\#\Delta$ ” if and only if \mathcal{T}_d accepts the input Δ .

To show that Algorithm 2 cannot be countered, consider any R.E. language to be chosen by Player “ \neq ”. This language, L_{\neq} , necessarily corresponds to the output of \mathcal{T}_x for some (finite) x . In total, Player “ $=$ ” can lose at most x rounds. In every subsequent round, its output will be identical to that of \mathcal{T}_x , and therefore identical to the bit chosen by Player “ \neq ”. \square

We see, therefore, that the complement of Algorithm 2’s language cannot be learned by any R.E. language. Player “ \neq ” cannot hope to win more than a finite number of rounds.

Note that these results do not necessitate that $\Sigma = \Sigma_1^0$, the R.E. languages. As long as Σ is rich enough to allow implementing Algorithm 2, the results hold. This is true, for example, for Σ sets that allow Oracle calls. In particular:

Corollary 3.1. *For all $i > 0$, Σ_i^0 is adversarially learnable by Σ_i^0 but not by Π_i^0 ; Π_i^0 is adversarially learnable by Π_i^0 but not by Σ_i^0 .*

Proof. To show the learnability results, we use Algorithm 2. To show the non-learnability results, we appeal to the symmetric nature of the game: if Player “ $=$ ” has a winning learning strategy, Player “ \neq ” does not. \square

5. Conventional learnability

To adapt the IMP game for the study of conventional (i.e., non-adversarial) learning and approximation, we introduce the notion of **nonadaptive strategies**.

Definition 3. A **nonadaptive strategy** is a language, L , over $\{0, 1\}^*$ such that

$$\forall u, v, |u| = |v| \Rightarrow (u \in L \Leftrightarrow v \in L),$$

where $|u|$ is the bit length of u .

Respective to an arbitrarily chosen (computable) enumeration w_1, w_2, \dots over the complete language, we define the function $NA()$ such that, for any language L , $NA(L)$ is the language such that

$$x \in NA(L) \Leftrightarrow w_{|x|} \in L.$$

Furthermore, for any collection of languages, Σ , we define $NA(\Sigma) = \{NA(L) | L \in \Sigma\}$. $NA(\Sigma)$ is the **nonadaptive application** of Σ .

To elucidate this definition, consider once again a (computable) enumeration, w_1, w_2, \dots over the complete language.

In previous sections, we have analysed the case where the two competing strategies are adaptive (i.e., general). This was the case of adversarial learning. Modelling the conventional learning problem is simply done by restricting Σ_{\neq} to nonadaptive strategies. The question of whether a strategy $L_{=}$ (or $D_{=}$) can learn L is the question of whether it can learn adversarially $NA(L)$. The reason this is so is because the bit output at any round i by a nonadaptive strategy is independent of any response made by either player at any previous round: at each round i , $O_{\neq}(i+1)$, the response of Player “ \neq ”, as defined in (5), is a function of Δ_i , a word composed of exactly i bits. Definition 3 now adds to this the restriction that the response must be invariant to the value of these i bits and must depend only on the bit length, i , which is to say on the round number. Regardless of what the strategy of Player “ $=$ ” is, the sequence $O_{\neq}(1), O_{\neq}(2), \dots$ output by Player “ \neq ” will always remain the same. Thus, a nonadaptive strategy for Player “ \neq ” is one where the player’s output is a predetermined, fixed string of bits, and it is this string that the opposing strategy of Player “ $=$ ” must learn to mimic.

Note, furthermore, that if Σ_{NA} is the set of all nonadaptive languages, then for every $i > 0$ we have

$$NA(\Sigma_i^0) = \Sigma_i^0 \cap \Sigma_{NA}. \quad (10)$$

The equality stems from the fact that calculating $w_{|x|}$ from x and vice versa (finding any x that matches $w_{|x|}$) is, by definition, recursive, so there is a reduction from any L to $NA(L)$ and back. If a language can be computed over the input $w_{|x|}$ by means of a certain nonempty set of quantifiers, no additional unbounded quantifiers are needed to compute it from x .

This leads us to Definition 4.

Definition 4. We say that a collection of languages Σ_{\neq} is (conventionally) **learnable** by a collection of strategies $\Sigma_{=}$ if $\minmax(\Sigma_{=}, NA(\Sigma_{\neq})) = 0$.

If a collection is learnable by Σ_1^0 , we simply say that it is **learnable**.

Corollary 3.2. For all $i > 0$, Σ_i^0 is learnable by Σ_i^0 . In particular, Σ_1^0 is learnable.

Proof. We have already shown (Corollary 3.1) that Σ_i^0 is adversarially learnable by Σ_i^0 , and $NA(\Sigma_i^0)$ is a subset of Σ_i^0 , as demonstrated by (10). \square

Constraining Player “ \neq ” to only be able to choose nonadaptive strategies can only lower the minmax value. Because it is already at 0, it makes no change: we are weakening the player that is already weaker. It is more rewarding to constrain Player “ $=$ ” and to consider the game $\text{IMP}(NA(\Sigma_i^0), \Sigma_i^0)$. Note, however, that this is equivalent to the game $\text{IMP}(\Sigma_i^0, NA(\Pi_i^0))$ under role reversal.

Theorem 4. Π_1^0 is learnable.

Proof. To begin, let us consider a simpler scenario than was discussed so far. Specifically, we will consider a scenario in which the feedback available to the learning algorithm at each point is not only Δ_n , the information of which rounds it had “won” and which it had “lost”, but also $O_=(n)$ and $O_\neq(n)$, what the bit output by each machine was, at every step.²

In this scenario, Player “=” can calculate a co-R.E. function by calculating its complement in round n and then reading the result as the complement to $O_=(n)$, which is given to it in all later rounds.

For example, at round n Player “=” may simulate a particular Turing machine, \mathcal{T} , in order to test whether it halts. If it does halt, the player halts and accepts the input, but it may also continue indefinitely. The end effect is that if \mathcal{T} halts then $O_=(n) = 1$ and otherwise it is 0. At round $n + 1$, Player “=” gets new inputs. (Recall that if one views the player as a Turing machine, it is effectively restarted at each round.) The new input in the real IMP game is Δ_n , but for the moment we are assuming a simpler version where the input is the pair of strings $(O_=(1) \dots O_=(n), O_\neq(1) \dots O_\neq(n))$. This being the case, though whether \mathcal{T} halts or not is in general not computable by a Σ_1^0 player, once a simulation of the type described here is run at round n , starting with round $n + 1$ the answer is available to the player in the form of $O_=(n)$, which forms part of its input.

More concretely, one algorithm employable by Player “=” against a known nonadaptive language $NA(L_\neq)$ is one that calculates “ $w_{2n+1} \notin L_\neq$?” (which is an R.E. function) in every $2n$ ’th round, and then uses this information in the next round in order to make the correct prediction. This guarantees $S(L_=(, NA(L_\neq)) \leq 1/2$. However, it is possible to do better.

To demonstrate how, consider that Player “=” can determine the answer to the question “ $|\{w_i, \dots, w_j\} \setminus L_\neq| \geq k$?” for any chosen i, j and k . The way to do this is to simulate simultaneously all $j + 1 - i$ Turing machine runs that calculate “ $w_l \notin L_\neq$?” for each $i \leq l \leq j$ and to halt if k of them halt. As with the previous example, by performing this algorithm at any stage n , the algorithm will then be able to read out the result as $O_=(n)$ in all later rounds.

Consider, now, that this ability can be used to determine $|\{w_i, \dots, w_j\} \setminus L_\neq|$ exactly (rather than simply bounding it) by means of a binary search, starting with the question “ $|\{w_i, \dots, w_j\} \setminus L_\neq| \geq 2^{m-1}$?” in the first round, and proceeding to increasingly finer determination of the actual set size on each later round. Player “=” can therefore determine the number of “1” bits in a set of $j + 1 - i = 2^m - 1$ outputs of a co-R.E. function in this way in only m queries, after which the number will be written in binary form, from most significant bit to least significant bit, in its $O_=(n)$ input. Once this cardinality has been determined, Player “=” can compute via a terminating computation the value of each of “ $w_l \in L_\neq$?”: the player will simulate, in parallel, all $j + 1 - i$ machines, and will terminate the computation either when the desired bit value is found via a halting of the corresponding machine, or until the full cardinality of halting machines has been reached,

²Because $O_=(n) \oplus O_\neq(n) \oplus \delta_n = 0$, using any two of these as input to the TM is equivalent to using all three, because the third can always be calculated from the others.

at which point, if the desired bit is not among the machines that halted, then the player can safely conclude that its computation will never halt.

Let $\{m_t\}$ be an arbitrary (computable) sequence with $\lim_{t \rightarrow \infty} m_t = \infty$. If Player “=” repeatedly uses m_t bits (each time picking the next value in the sequence) of its own output in order to determine Player “ \neq ”’s next $2^{m_t} - 1$ bits, the proportion of bits determined correctly by this will approach 1.

However, the actual problem at hand is one where Player “=” does not have access to its own output bits, $(O_=(1), \dots, O_=(n))$. Rather, it can only see $(\delta_1, \dots, \delta_n)$, the exclusive or (xor) values of its bits and those of Player “ \neq ”. To deal with this situation, we use a variation over the strategy described above.

First, for convenience, assume that Player “=” knows the first m_0 bits to be output by Player “ \neq ”. Knowing Player “ \neq ”’s bits and having visibility as to whether they are the same or different to Player “=”’s bits give, together, Player “=” access to its own past bits.

Now, it can use these first m_0 bits in order to encode, as before, the cardinality of the next $2^{m_0} - 1$ bits, and by this also their individual values (as was demonstrated previously with the calculation of “ $w_l \in L_{\neq}?$ ”). This now gives Player “=” the ability to win every one of the next $2^{m_0} - 1$ rounds. However, instead of utilising this ability to the limit, Player “=” will only choose to win the next $2^{m_0} - 1 - m_1$, leaving the remaining m_1 bits free to be used for encoding the cardinality of the next $2^{m_1} - 1$. This strategy can be continued to all m_t . The full list of criteria required of the sequence $\{m_t\}$ for this construction to work and to ultimately lead to $S(L_=(, NA(L_{\neq})) = 0$ is:

1. $\lim_{t \rightarrow \infty} m_t = \infty$.
2. $\forall t, m_{t+1} \leq 2^{m_t} - 1$.
3. $\lim_{t \rightarrow \infty} \frac{m_{t+1}}{2^{m_t}} = 0$.

A sequence satisfying all these criteria can easily be found, e.g. $m_t = t + 2$.

Two problems remain to be solved: (1) How to determine the value of the first m_0 bits, and (2) how to deal with the fact that L_{\neq} is not known.

We begin by tackling the second of these problems. Because L_{\neq} is not known, we utilise a strategy of enumerating over the possible languages, similar to what is done in Algorithm 2. That is to say, we begin by assuming that $\text{co-}L_{\neq} = L_0$ and respond accordingly. Then, if we detect that the responses from Player “ \neq ” do not match those of L_0 we progress to assume that $\text{co-}L_{\neq} = L_1$, etc.. We are not always in a position to tell if our current hypothesis of L_{\neq} is correct, but we can verify that it matches at least the first $2^{m_t} - m_{t+1} - 1$ bits of each $2^{m_t} - 1$ set. If Player “=” makes any incorrect predictions during any of these $2^{m_t} - m_{t+1} - 1$ rounds, it can progress to the next hypothesis. We note that it is true that Player “=” can remain mistaken about the identity of L_{\neq} forever, as long as L_{\neq} is such that the first $2^{m_t} - m_{t+1} - 1$ predictions of every $2^{m_t} - 1$ are correct, but because these correct predictions alone are enough to ensure $S(L_=(, NA(L_{\neq})) = 0$, the question of whether the correct L_{\neq} is ultimately found or not is moot.

To tackle the remaining problem, that of determining m_0 bits of $L_=($ in order to bootstrap the process, we make use of mixed strategies.

Consider a mixed strategy involving probability $1/2^{m_0}$ for each of 2^{m_0} strategies, differing only by the m_0 bits they assign as the first bits for each language in order to bootstrap the learning process. If $\text{co-}L_{\neq} = L_0$, of the 2^{m_0} strategies one will make the correct guess regarding the first m_0 input bits, after which that strategy can ensure $S(L_{=}, NA(L_{\neq})) = 0$. However, note that, if implemented as described so far, this is not the case for any other L_i . Suppose, for example, that $\text{co-}L_{\neq} = L_1$. All 2^{m_0} strategies begin by assuming, falsely, that $\text{co-}L_{\neq} = L_0$, and all may discover later on that this assumption is incorrect, but they may do so at different rounds. Because of this, a counter-strategy can be designed to fool all 2^{m_0} learner strategies.

To avoid this pitfall, all strategies must use the same bit positions in order to bootstrap learning for each L_i , so these bit positions must be pre-allocated. We will use bits $a^2, \dots, a^2 + m_0 - 1$ in order to bootstrap the i 'th hypothesis, for some known $a = a(i)$, regardless of whether the hypothesis $L_{=} = L_i$ is known to require checking before these rounds, after, or not at all. The full set of rounds pre-allocated in this way still has only density zero among the integers, so even without a win for Player “=” in any of these rounds its final payoff remains 1.

Suppose, now, that L_i is still not the assumption currently being verified (or falsified) at rounds $a^2, \dots, a^2 + m_0 - 1$. The Hamming weight (number of “1”s) of which $2^{m_0} - 1$ bits should be encoded by Player “=” in these rounds’ bits? To solve this, we will pre-allocate to each hypothesis an infinite number of bit positions, which, altogether for all hypotheses, still amount to a set of density 0 among the integers. The hypothesis will continuously predict the values of this pre-allocated infinite sequence of bits until it becomes the “active” assumption. If and when it does, it will expand its predictions to all remaining bit positions.

This combination of 2^{m_0} strategies, of which one guarantees a payoff of 1, therefore guarantees in total an expected payoff of at least $1/2^{m_0}$. We want to show, however, that $\min_{\max}(\Sigma_1^0, NA(\Pi_1^0)) = 0$. To raise from $1/2^{m_0}$ to 1, we describe a sequence of mixed strategies for which the expected payoff for Player “=” converges to 1.

The k 'th element in the sequence of mixed strategies will be composed of $2^{m_0 k}$ equal probability pure strategies. The strategies will follow the algorithm so far, but instead of moving from the hypothesis $\text{co-}L_{\neq} = L_i$ to $\text{co-}L_{\neq} = L_{i+1}$ after a single failed attempt (which may be due to incorrect bootstrap bits), the algorithm will try each L_i language k times. In total, it will guess at most $m_0 k$ bits for each language, which are the $m_0 k$ bits defining the strategy.

This strategy ensures a payoff of at least $1 - (1 - 1/2^{m_0})^k$, so converges to 1, as desired, for an asymptotically large k .

The full algorithm is described in Algorithm 3. It uses the function *triangle*, defined as follows: let

$$\text{base}(x) = \left\lfloor \frac{\lfloor \sqrt{8x+1} \rfloor - 1}{2} \right\rfloor$$

and

$$\text{triangle}(x) = x - \text{base}(x)(\text{base}(x) + 1)/2. \quad (11)$$

The value of $triangle(x)$ for $x = 0, 1, 2, \dots$ equals

$$0, 0, 1, 0, 1, 2, 0, 1, 2, 3, 0, 1, 2, 3, 4, 0, 1, 2, 3, 4, 5, \dots,$$

describing a triangular walk through the nonnegative integers.

The algorithm is divided into two stages. In Step 1, the algorithm simulates its actions in all previous rounds, but without simulating any (potentially non-halting) Turing machine associated with any hypothesis. The purpose of this step is to determine which hypothesis (choice of Turing machine and bootstrapping) is to be used for predicting the next bit. Once the hypothesis is determined, Step 2 once again simulates all previous rounds, only this time simulating the chosen hypothesis wherever it is the active hypothesis. In this way, the next bit predicted by the hypothesis can be determined.

The specific $\{m_t\}$ sequence used in Algorithm 3 is $m_t = t + 2$ (which was previously mentioned as an example of a sequence satisfying all necessary criteria). □

Some corollaries follow immediately.

Corollary 4.1. *There exists a probabilistic Turing machine that is able to learn any language in Π_1^0 with probability 1.*

Proof. Instead of using a mixed strategy, it is possible to use probabilistic Turing machines in order to generate the m_0 guessed bits that bootstrap each hypothesis. In this case, there is neither a need for a mixed strategy nor a need to consider asymptotic limits: a single probabilistic Turing machine can perform a triangular walk over the hypotheses for L_{\neq} , investigating each option an unbounded number of times. The probability that for the correct L_{\neq} at least one bootstrap guess will be correct in this way equals 1.

The method for doing this is essentially the same as was described before. The only caveat is that because the probabilistic TM is re-initialised at each round and because it needs, as part of the algorithm, to simulate its actions in all previous rounds, the TM must have a way to store its random choices, so as to make them accessible in all later rounds.

The way to do this is to extend the hypothesis “bootstrap” phase from m_0 bits to $2m_0$ bits. In each of the first m_0 bits, the TM outputs a uniform random bit. The δ_n bit available to it in all future rounds is then this random bit xor the output of Player “ \neq ”. δ_n is therefore also a uniform random bit. In this way, in all future rounds the TM has access to these m_0 consistent random bits. It can then use these in the second set of m_0 bootstrap bits as was done with the j value in the deterministic set-up. □

We note, as before, that the construction described continues to hold, and therefore the results remain true, even if Oracles are allowed, that are accessible to both players, and, in particular, the results hold for any Π_i^0 with $i > 0$:

Corollary 4.2. *For all $i > 0$, Π_i^0 is learnable by Σ_i^0 .*

Furthermore:

Algorithm 3 Algorithm for learning any co-R.E. language

```

1:                                     ▷ The strategy is a uniform mixture of  $4^k$  algorithms.
2:                                     ▷ We describe the  $j$ 'th algorithm.
3: function CALCULATE BIT( $\Delta$ )
4:    $n \leftarrow$  length of  $\Delta$                                      ▷ The round number. Let  $\Delta = \delta_1, \dots, \delta_n$ .
5:                                     ▷ Step 1: Identify  $h$ , the current hypothesis.
6:    $NonActiveHypotheses \leftarrow \{\}$ 
7:    $PredPos \leftarrow \{\}$                                      ▷ A set managing which positions are predicted by which
   hypothesis.
8:   for  $i \in 0, \dots, n$  do
9:     if  $\exists(h, S, S') \in PredPos$  such that  $i \in S$  then
10:      Let  $h, S, S'$  be as above.
11:                                     ▷  $h$  = hypothesis number.
12:                                     ▷  $S$  = predicted positions.
13:                                     ▷  $S'$  = next positions to be predicted.
14:      Let  $m$  be such that  $2^{m-1} - 1 = |S'|$ . ▷ We only construct  $S'$  that have such
   an  $m$ .
15:    else if  $\exists a, h$  such that  $a^2 = i$ ,  $h = triangle(a)$  and  $h \notin NonActiveHypotheses$ 
   then
16:      ▷ First bootstrap bit for hypothesis  $h$ .
17:      Let  $h$  be as above.
18:       $S \leftarrow \{\}$ 
19:       $S' \leftarrow \{i, i + 1\}$ 
20:       $bootstrap(h) \leftarrow i$ 
21:       $m \leftarrow 2$ 
22:    else if  $i = n$  then                                     ▷ Unusable bits.
23:      Accept input.                                           ▷ Arbitrary choice.
24:    else
25:      Next  $i$ .
26:    end if
27:     $e \leftarrow |\{x \in S | x > i\}|$ 
28:    if  $e \geq m$  then
29:      ▷ These bits are predicted accurately for the correct hypothesis.
30:      if  $i < n$  and  $\delta_{i+1} = 1$  then
31:        ▷ Incorrect prediction, so hypothesis is false.
32:         $NonActiveHypotheses \leftarrow NonActiveHypotheses \cup \{h\}$ 
33:         $PredPos \leftarrow \{(\tilde{h}, \tilde{S}, \tilde{S}') \in PredPos | \tilde{h} \neq h\}$ 
34:      end if
35:    else if  $e = m - 1$  then ▷ Bits with  $e < m$  are used to encode next bit counts.
36:       $\tilde{S} \leftarrow \{\}$                                      ▷ New positions to predict on.
37:       $p \leftarrow \max(S')$ 

```

```

38:         while  $|\tilde{S}| < 2^m - 1$  do
39:              $p \leftarrow p + 1$ 
40:             if  $(\exists a, b$  such that  $b \in \{0, 1\}$ ,  $a^2 + b = p$  and  $h = \text{triangle}(a)$ ) or  $(h =$ 
                $\text{mex}(\text{NonActiveHypotheses})$  and  $\nexists a, b, \tilde{h}$  such that  $b \in \{0, 1\}$ ,  $a^2 + b = p$ ,  $\tilde{h} = \text{triangle}(a)$ ,
                $\tilde{h} \notin \text{NonActiveHypotheses})$  then
41:                  $\triangleright$  “ $\text{mex}(T)$ ” is the minimum nonnegative integer not appearing in  $T$ .
42:                  $\tilde{S} \leftarrow \tilde{S} \cup \{p\}$ 
43:             end if
44:         end while
45:          $\text{PredPos} \leftarrow \text{PredPos} \cup (h, S', \tilde{S})$ 
46:     end if
47: end for
48:                                      $\triangleright$  Step 2: Predict, assuming  $h$ .
49:      $i \leftarrow \text{bootstrap}(h)$ 
50:      $S \leftarrow \{i, i + 1\}$ 
51:      $M \leftarrow h \text{ div } k$                                       $\triangleright T_M$  is the machine to be simulated.  $x \text{ div } y \stackrel{\text{def}}{=} \lfloor x/y \rfloor$ .
52:      $\text{try} \leftarrow h \text{ mod } k$                                       $\triangleright$  The try number of this machine.
53:      $\text{Prediction}(i) \leftarrow (j \text{ div } 4^{\text{try}}) \text{ mod } 2$ 
54:      $\text{Prediction}(i + 1) \leftarrow (j \text{ div } (2 \cdot 4^{\text{try}})) \text{ mod } 2$ 
55:     for  $i \in 0, \dots, n$  do
56:         if  $\exists S, S', (h, S, S') \in \text{PredPos}$  and  $i \in S$  then
57:             Let  $m$  be such that  $2^m - 1 = |S'|$ .
58:              $e \leftarrow |\{x \in S \mid x > i\}|$ 
59:             if  $e = m - 1$  then
60:                  $\text{counter} \leftarrow 0$                                       $\triangleright$  Number of 1's in  $S'$ .
61:             end if
62:             if  $e \geq m$  then
63:                 if  $i = n$  then
64:                     if  $\text{Prediction}(i) = 1$  then
65:                         Accept input.
66:                     else
67:                         Reject input.
68:                     end if
69:                 end if
70:             else if  $i = n$  then
71:                 Simulate  $T_M$  simultaneously on all inputs in  $S'$  until  $\text{counter} + 2^e$  are
               accepted.
72:                  $\triangleright$  If this simulation does not terminate, this is a rejection of the input.
73:                 Accept input.
74:             else

```

```

75:      if  $Prediction(i) \neq \delta_i$  then                                 $\triangleright$  Previous simulation terminated.
76:           $counter \leftarrow counter + 2^e$                                  $\triangleright$  Binary search.
77:      end if
78:      if  $e = 0$  then                                 $\triangleright$   $counter$  holds the number of terminations in  $S'$ .
79:          Simulate  $T_M$  simultaneously on all inputs in  $S'$  until  $counter$  are ac-
          cepted.                                                 $\triangleright$  Guaranteed to halt, if hypothesis is
          correct.
80:          Let  $Prediction(x)$  be 0 on all  $x \in S'$  that terminated, 1 otherwise.
81:      end if
82:  end if
83: end if
84: end for
85: end function

```

Corollary 4.3. *For all $i > 0$, the collection of languages learnable by Σ_i^0 is a strict superset of $\Sigma_i^0 \cup \Pi_i^0$.*

Proof. We have already shown that Σ_i^0 and Π_i^0 are both learnable by Σ_i^0 . Adding the Σ_i^0 languages as additional hypotheses to Algorithm 3 we can see that the set $\Sigma_i^0 \cup \Pi_i^0$ is also learnable.

To give one example of a family of languages beyond this set which is also learnable by Σ_i^0 , consider the following. Let $\Sigma_i^{(c)}$, for a fixed $c > 1$, be the set of languages recognisable by a Δ_0^0 Turing machine which can make at most c calls to a Σ_i^0 Oracle.

This set contains Σ_i^0 and Π_i^0 , but it also contains, for example, the xor of any two languages in Σ_i^0 , which is outside of $\Sigma_i^0 \cup \Pi_i^0$, and therefore strictly beyond the i 'th level of the arithmetic hierarchy.

We will adapt Algorithm 3 to learn $\Sigma_i^{(c)}$. The core of Algorithm 3 is its ability to use m bits of Δ_n in order to predict $2^m - 1$ bits. We will, instead, use cm bits in order to predict the same amount. Specifically, we will use the first m bits in order to predict the result of the first Oracle call in each of the predicted $2^m - 1$ positions, the next m bits in order to predict the second Oracle call in each of the predicted $2^m - 1$ positions, and so on.

In total, for this to work, all we need is to replace criterion 2 in our list of criteria for the $\{m_t\}$ sequence with the new criterion

$$\forall t, cm_{t+1} \leq 2^{m_t} - 1.$$

An example of such a sequence is $m_t = t + \max(c, 5)$. □

In fact, Algorithm 3 can be extended even beyond what was described in the proof to Corollary 4.3. For example, instead of using a constant c , it is possible to adapt the algorithm to languages that use $c(n)$ Oracle calls at the n 'th round, for a sufficiently low-complexity $c(n)$ by similar methods.

Altogether, it seems that R.E. learning is significantly more powerful than being able to learn merely the first level of the arithmetic hierarchy, but we do not know whether it

can learn every language in Δ_2^0 . Indeed, we have no theoretical result that implies R.E. learning cannot be even more powerful than the second level of the arithmetic hierarchy.

A follow-up question which may be asked at this point is whether it was necessary to use a mixed strategy, as was used in the proof of Theorem 4, or whether a pure strategy could have been designed to do the same.

In fact, no pure strategy would have sufficed:

Lemma 4.1. *For all i ,*

$$\inf_{L_= \in \Sigma_i^0} \sup_{L_\neq \in NA(\Pi_i^0)} S(L_-, L_\neq) = 1.$$

This result is most interesting in the context of Corollary 4.1, because it describes a concrete task that is accomplishable by a probabilistic Turing machine but not by a deterministic Turing machine.

Proof. We devise for each L_- a specific L_\neq antidote. The main difficulty in doing this is that we cannot choose, as before, $L_\neq = \text{co-}L_-$, because L_\neq is now restricted to be nonadaptive, whereas L_- is general.

However, consider L_\neq such that its bit for round k is the complement of L_- 's response on $\Delta_{k-1} = 1^{k-1}$. This is a nonadaptive strategy, but it ensures that Δ_k will be 1^k for every k . Effectively, L_\neq describes L_- 's “red herring sequence”. \square

6. Approximability

When *both* players' strategies are restricted to be nonadaptive, they have no means of learning each other's behaviours: determining whether their next output bit will be 0 or 1 is done solely based on the present round number, not on any previous outputs. The output of the game is therefore solely determined by the dissimilarity of the two independently-chosen output strings.

Definition 5. We say that a collection of languages Σ_\neq is **approximable** by a collection of strategies $\Sigma_=-$ if $\minmax(NA(\Sigma_-), NA(\Sigma_\neq)) = 0$.

If a collection is approximable by Σ_1^0 , we simply say that it is **approximable**.

In this context it is clear that for any Σ

$$\sup_{L_\neq \in NA(\Sigma)} \inf_{L_- \in NA(\Sigma)} S(L_-, L_\neq) = 0,$$

because L_- can always be chosen to equal L_\neq , but unlike in the case of adversarial learning, here mixed strategies do make a difference.

Though we do not know exactly what the value of $\minmax(NA(\Sigma_1^0))$ is, we do know the following.

Lemma 4.2. *If $D_ =$ and $D_ \neq$ are mixed strategies from $NA(\Sigma_1^0)$, then*

$$\sup_{D_ \neq} \inf_{D_ =} E \left(\limsup_{N \rightarrow \infty} \sum_{n=1}^N \frac{\delta_n}{N} \right) \geq \frac{1}{2} \quad (12)$$

and

$$\inf_{D_ =} \sup_{D_ \neq} E \left(\limsup_{N \rightarrow \infty} \sum_{n=1}^N \frac{\delta_n}{N} \right) \geq \frac{1}{2}, \quad (13)$$

where δ_n is as in the definition of the IMP game.

In other words, Player “ \neq ” can always at the very least break even, from a lim sup perspective.

Proof. Let $D_ \neq$ be a mixture of the following two strategies: all zeros (L_0), with probability 1/2; all ones (L_1), with probability 1/2. By the triangle inequality, we have that for any language $L_ =$,

$$E \left(\limsup_{N \rightarrow \infty} \sum_{n=1}^N \frac{\delta_n}{N} \right) = \frac{\text{DisSim}(L_ =, L_0) + \text{DisSim}(L_ =, L_1)}{2} \geq \frac{\text{DisSim}(L_0, L_1)}{2} = \frac{1}{2},$$

and because this is true for each $L_ =$ in $D_ =$, it is also true in expectation over all $D_ =$. The fact that $D_ \neq$ is independent of $D_ =$ in the construction means that this bound is applicable for both (12) and (13). \square

Just as interesting (and with tighter results) is the investigation of lim inf. We show

Lemma 4.3.

$$\inf_{L_ = \in NA(\Sigma_1^0)} \sup_{L_ \neq \in NA(\Sigma_1^0)} \liminf_{N \rightarrow \infty} \sum_{n=1}^N \frac{\delta_n}{N} = \sup_{L_ \neq \in NA(\Sigma_1^0)} \inf_{L_ = \in NA(\Sigma_1^0)} \liminf_{N \rightarrow \infty} \sum_{n=1}^N \frac{\delta_n}{N} = 0, \quad (14)$$

where δ_n is as in the definition of the IMP game.

Proof. Let $\text{triangle}(x)$ be as in (11), and let $\text{caf}(x)$ be the maximum integer, y , such that $y! \leq x$.

The language $L_ =$ will be defined by

$$w_i \in L_ = \Leftrightarrow w_i \in L_{\text{triangle}(\text{caf}(i))},$$

where L_0, L_1, L_2, \dots is an enumeration over all R.E. languages.

To prove that for any j , if $L_ \neq = L_j$ the claim holds, let us first join the rounds into “super-rounds”, this being the partition of the rounds set according to the value of $y = \text{caf}(i)$. At each super-round, $L_ =$ equals a specific L_x , and by the end of the super-round, a total of $(y-1)/y$ of the total rounds will have been rounds in which $L_ =$ equals this L_x . Hence, the Hamming distance between the two (the number of differences) at this time is at most $1/y$ of the string lengths. Because each choice of x repeats an infinite number of times, the lim inf of this proportion is 0. \square

With this lemma, we can now prove Theorem 1.

Proof. The theorem is a direct corollary of the proof of Lemma 4.3, because the complement of the language $L_=_$ that was constructed in the proof to attain the infimum can be used as \bar{L} . \square

Combining Lemma 4.2 and 4.3 with the definition of the payoff function in (1), we get, in total:

Corollary 4.4.

$$1/4 \leq \maxmin(NA(\Sigma_1^0)) \leq 1/2$$

and

$$1/4 \leq \minmax(NA(\Sigma_1^0)) \leq 1/2.$$

Though we have the exact value of neither maxmin nor minmax in this case, we do see that the case is somewhat unusual in that neither player has a decisive advantage.

7. Conclusions and further research

We have introduced the IMP game as an arena within which to test the ability of algorithms to learn and be learnt, and specifically investigated three scenarios:

Adversarial learning, where both algorithms are simultaneously trying to learn each other by observations.

Non-adversarial (conventional) learning, where an algorithm is trying to learn a language by examples.

Approximation, where languages (or language distributions) try to mimic each other without having any visibility to their opponent's actions.

In the case of adversarial learning, we have shown that Σ_i^0 can learn Σ_i^0 but not Π_i^0 .

In conventional learning, however, we have shown that Σ_i can learn Σ_i^0 , Π_i^0 and beyond into the $(i + 1)^{\text{th}}$ level of the arithmetic hierarchy, but this learnability is yet to be upper-bounded. Our conjecture is that the class of learnable languages is strictly a subset of Δ_2^0 . If so, then this defines a new class of languages between the first and second levels of the arithmetic hierarchy, and, indeed, between any consecutive levels of it.

Regarding approximability, we have shown that (unlike in the previous results) no side has the absolute upper hand in the game, with the game value for Player “ \neq ”, if it exists, lying somewhere between $1/4$ and $1/2$. We do not know, however, whether the game is completely unbiased or not.

An investigation of adversarial learning in the context of recursive languages was given as a demonstration of the fact that in IMP it may be the case that no Nash equilibrium exists at all, and pure-strategy learning was given as a concrete example of a task where probabilistic Turing machines have a provable advantage over deterministic ones.

References

- [1] D.L. Dowe. Foreword re C. S. Wallace. *Computer Journal*, 51(5):523–560, September 2008. Christopher Stewart WALLACE (1933-2004) memorial special issue.
- [2] D.L. Dowe. Minimum Message Length and statistically consistent invariant (objective?) Bayesian probabilistic inference – from (medical) “evidence”. *Social Epistemology*, 22(4):433–460, Oct–Dec 2008.
- [3] D.L. Dowe. MML, hybrid Bayesian network graphical models, statistical consistency, invariance and uniqueness. In Bandyopadhyay, P.S. and Forster, M.R., editor, *Handbook of the Philosophy of Science – Volume 7: Philosophy of Statistics*, pages 901–982. Elsevier, 2011.
- [4] D.L. Dowe. Introduction to Ray Solomonoff 85th Memorial Conference. In *Proceedings of Solomonoff 85th memorial conference – Lecture Notes in Artificial Intelligence (LNAI)*, volume 7070, pages 1–36. Springer, 2013.
- [5] D.L. Dowe, J. Hernández-Orallo, and P.K. Das. Compression and intelligence: Social environments and communication. In *AGI: 4th Conference on Artificial General Intelligence – Lecture Notes in Artificial Intelligence (LNAI)*, pages 204–211, 2011.
- [6] G.W. Flake. *The Computational Beauty of Nature: Computer Explorations of Fractals, Chaos, Complex Systems, and Adaptation*. A Bradford book. Cambridge, Massachusetts, 1998.
- [7] E.M. Gold. Language identification in the limit. *Information and Control*, 10(5):447–474, 1967.
- [8] J. Hernández-Orallo, D.L. Dowe, S. España-Cubillo, M.V. Hernández-Lloreda, and J. Insa-Cabrera. On more realistic environment distributions for defining, evaluating and developing intelligence. In *AGI: 4th Conference on Artificial General Intelligence – Lecture Notes in Artificial Intelligence (LNAI)*, volume 6830, pages 82–91. Springer, 2011.
- [9] Ling Huang, Anthony D. Joseph, Blaine Nelson, Benjamin I.P. Rubinstein, and J. D. Tygar. Adversarial machine learning. In *Proceedings of the 4th ACM Workshop on Security and Artificial Intelligence, AISec ’11*, pages 43–58, New York, NY, USA, 2011. ACM.
- [10] D.K. Lewis and J.S. Richardson. Scriven on human unpredictability. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 17(5):69–74, October 1966.
- [11] Wei Liu and Sanjay Chawla. A Game Theoretical Model for Adversarial Learning. In Saygin, Y and Yu, JX and Kargupta, H and Wang, W and Ranka, S and Yu,

- PS and Wu, XD, editor, *2009 IEEE INTERNATIONAL CONFERENCE ON DATA MINING WORKSHOPS (ICDMW 2009)*, pages 25–30. Knime; Mitre; CRC Press, 2009. 9th IEEE International Conference on Data Mining, Miami Beach, FL, DEC 06-09, 2009.
- [12] Daniel Lowd and Christopher Meek. Adversarial learning. In *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, KDD '05, pages 641–647, New York, NY, USA, 2005. ACM.
 - [13] J. Nash. Non-cooperative Games. *The Annals of Mathematics*, 54(2):286–295, 1951.
 - [14] J.v. Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.
 - [15] Hartley Rogers, Jr. *Theory of Recursive Functions and Effective Computability*. MIT Press, Cambridge, MA, second edition, 1987.
 - [16] M. Scriven. An essential unpredictability in human behavior. In B.B. Wolman and E. Nagel, editors, *Scientific Psychology: Principles and Approaches*, pages 411–425. Basic Books (Perseus Books), 1965.
 - [17] R.J. Solomonoff. Complexity-based induction systems: Comparisons and convergence theorems. *IEEE Transaction on Information Theory*, IT-24(4):422–432, 1978.
 - [18] R.J. Solomonoff. Algorithmic probability: Theory and applications. In F. Emmert-Streib and M. Dehmer, editors, *Information Theory and Statistical Learning, Springer Science and Business Media*, pages 1–23. Springer, N.Y., U.S.A., 2009.
 - [19] R.J. Solomonoff. Algorithmic probability, heuristic programming and AGI. In *Proceedings of the Third Conference on Artificial General Intelligence, AGI 2010*, pages 251–257, Lugano, Switzerland, March 2010. IDSIA.
 - [20] R.J. Solomonoff. Algorithmic probability – its discovery – its properties and application to strong AI. In H. Zenil, editor, *Randomness Through Computation: Some Answers, More Questions*, pages 1–23. World Scientific Publishing Co., Inc., River Edge, NJ, USA, 2011.
 - [21] A.M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc.*, 42:230–265, 1936.
 - [22] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.